arXiv:1512.08212v2 [cs.CV] 22 May 2016

Improving Facial Analysis and Performance Driven Animation through Disentangling Identity and Expression

David Rim*[a], Sina Honari*[b], Md Kamrul Hasan[a], Chris Pal[a]

[a]*Département de génie informatique et génie logiciel*
École Polytechnique Montréal
Montréal, Québec, Canada, H3T 1J4
[b]*Département d'informatique et de recherche opérationnelle*
Université de Montréal
Montréal, Québec, Canada, H3C 3J7

Abstract

We present techniques for improving performance driven facial animation, emotion recognition, and facial key-point or landmark prediction using learned identity invariant representations. Established approaches to these problems can work well if sufficient examples and labels for a particular identity are available and factors of variation are highly controlled. However, labeled examples of facial expressions, emotions and key-points for new individuals are difficult and costly to obtain. In this paper we improve the ability of techniques to generalize to new and unseen individuals by explicitly modeling previously seen variations related to identity and expression. We use a weakly-supervised approach in which identity labels are used to learn the different factors of variation linked to identity separately from factors related to expression. We show how probabilistic

Email addresses: daverim@gmail.com (David Rim*), sina.honari@umontreal.ca (Sina Honari*), md-kamrul.hasan@polymtl.ca (Md Kamrul Hasan), christopher.pal@polymtl.ca (Chris Pal)

Published in Image and Vision Computing Journal (IMAVIS)

modeling of these sources of variation allows one to learn identity-invariant representations for expressions which can then be used to identity-normalize various procedures for facial expression analysis and animation control. We also show how to extend the widely used techniques of active appearance models and constrained local models through replacing the underlying point distribution models which are typically constructed using principal component analysis with identity-expression factorized representations. We present a wide variety of experiments in which we consistently improve performance on emotion recognition, markerless performance-driven facial animation and facial key-point tracking.
Keywords: Factorization techniques; Emotion recognition; Graphical Models; Performance driven animation; Facial expression analysis

1. Introduction

One of the primary sources of variation in facial images is identity. Although this is an obvious statement, many approaches to vision tasks other than facial recognition do not directly account for the interaction between identity-related variation and other sources. However, many facial image datasets are subdivided by subject identity and this provides additional information that is often unused. This paper deals with the natural question of how to effectively use identity information in order to improve tasks other than identity recognition. In particular, our primary motivations, applications of interest and goals are to develop methods for facial expression analysis and performance driven facial animation that are less identity dependant.

*Joint first authors.

Recently there has been work on facial recognition in which identity is separated from other sources of variation in 2D image data in a fully probabilistic way [37]. In this model, the factors are assumed to be additive and independent. This procedure can be interpreted as a probabilistic version of Canonical Correlation Analysis (CCA) presented by Bach [2], or as a standard factor analysis with a particular structure in the factors. In this paper, we investigate and extend the use of this probabilistic approach to separate sources of variation, but unlike prior work which focuses on inferences about identity, we focus on facial expression analysis and facial animation tasks. This includes performance-driven animation, emotion recognition, and key-point tracking. Our goal is to use learned representations so as to create automated techniques for expression analysis and animation that better generalize across identities. We show here how disentangling factors of variation related to identity can indeed yield improved results. In many cases we show how to use the learned representations as input to discriminative classification methods.

In our experimental work, we apply this learning technique to a wide variety of different input types including: raw pixels, key-points and the pixels of warped images. We evaluate our approach by predicting: standard emotion labels, facial action units, and 'bone' positions or animation sliders which are widely used in computer animation. We go on to show how it is indeed possible to improve the facial keypoint prediction performance of active appearance models (AAMs) on unseen identities as well as increase the performance of constrained local models (CLMs) through our identity-expression factorization extensions to these widely used techniques. Our evaluation tasks and a comparison of the types of data being used as input for each experiment are summarized in Table 1.

The rest of this manuscript is structured as follows: In Section 2, we discuss

the facial expression analysis, performance driven animation and keypoint tracking applications that serve as the ultimate goals of our work in more detail and review relevant previous work. The various methods we present here build in particular on the work of Prince et al. [37] in which a linear Gaussian probabilistic model was proposed to explicitly separate factors of variation due to identity versus expression. In section 3, we present this model as a way to disentangle factors of variation arising from identity and expression variation. While Prince et al. used this type of identity-expression analysis to make inferences about identity, our work here focuses on how disentangling such factors can be used to make inferences about facial expression. Indeed, as discussed above, facial expression analysis and the applications of detailed facial expression analysis to computer animation is the motivating goal of our work here.

The first set of contributions of our work are presented in section 3.2. These contributions consist of a wide variety of novel techniques for using learned identity-expression representations for common goals related to expression analysis and computer animation. We provide novel techniques and experiments in which we predict: emotion labels, facial action units, facial keypoints, and animation control points as summarized in Table 1. Of particular note is the fact that we propose a novel formulation for identity normalizing facial images which we use for facial action unit recognition, emotion recognition and animation control. We find that using this identity-normalized representation leads to improved results across this wide variety of the expression analysis tasks.

We also go on to extend the underlying identity and expression analysis technique in two important directions, providing two additional technical contributions. First, in Section 4 we show that identity-expression analysis can be used

in place of the principal component analysis (PCA) technique that is widely used in active appearance models (AAMs). This procedure yields what we call an identity-expression factorized AAM, or IE-AAM. We present the modifications that are necessary to integrate this approach into the AAM framework of [34]. Our experiments show that IE-AAMs can increase the performance of PCA-AAMs dramatically when no training data is available for a given subject. We also found that the use of IE-AAMs eliminated the convergence errors observed with PCA-AAMs. Secondly, in Section 5 we show that constrained local models (CLMs) can also be reformulated, extended and improved through using an underlying identity-expression analysis model. Our reformulation also provides a novel energy function and minimization formulation for CLMs in general.

2. Our Applications of Interest and Relevant Prior Work

Performance-driven facial animation. Performance-driven animation is the task of controlling a facial animation via images of a performer. This is a common process in the entertainment industry. In many cases, this problem is generally handled by the placement of special markers on the performer [52]. However, here we are interested in developing markerless motion capture techniques in which we employ machine learning methods and minimize manual intervention.

Many marker-less facial expression analysis methods rely either on tracking points using optical flow [16] or fitting Active Appearance Models [30]. Morphable models in [6], [36], [40] have also been investigated. More recently, 3D data and reconstruction is used to fit directly to the performer [51], [54]. These methods, however, often require additional data, for example, multiview stereo [53], [4] or structured light [7], [51]. Sandbach *et al.* provide a thorough review of

3D expression recognition techniques in [38]. In the end, these methods usually work by providing dense correspondences which then require a re-targeting step.

However, a simpler and often used approach in industry does not rely on markers and simply uses the input video of a facial performance. The idea is to use a direct 2D to 3D mapping based on regressing image features [21] to 3D model parameters. This method works well but is insufficiently automatic. Each video is mapped to 3D model parameters, possibly with interpolation between frames. Key-point based representations often require both data and training time that compares unfavorably to this simpler approach given the re-targeting step. The primary benefit of key-point based representations appears to be a degree of natural identity-invariance. We provide experiments on performance driven animation, in which we predict bone positions using the well known Japanese Female Facial Expression (JAFFE) Database [33] as input as well as an experiment using professional helmet camera based video used for high quality, real world animations.

Emotion recognition. There is a large amount of overlap in the objectives for performance-driven facial animation and the objectives of detailed emotion recognition. Automatic emotion recognition has largely focused on the facial action coding system or FACS [15] as an auxiliary task [46]. In our work here we shall predict action unit (AU) values that have been annotated for the well known extended Cohn Kanade database [32]. In some cases, a FACs representation may be a useful intermediate representation for animation control.

Emotions can also be encoded via the simpler and more intuitive labels for well known emotion types of "Anger," "Disgust," "Fear," "Happiness," "Sadness," "Surprise". We evaluate the performance of recognizing the annotations of these emotions for JAFFE as well as CK+ in which "Contempt" is also included. These

labels correspond to Paul Ekman's famous set of six basic emotions plus "Contempt" which is a part of his extended list of emotions.

Many existing approaches for emotion recognition also require significant preprocessing for contrast normalization and alignment, which helps to alleviate pose variation [43]. Indeed, an intermediary task that has often been used as a preprocessing phase to produce emotion recognition systems is key point detection and tracking [42], [3], [45], [9]. Tracked key-points themselves can also be used to control animations directly or indirectly via warping operations on images as we explore in our experiments on professional facial performance video obtained via helmet cameras.

Key-point localization. Active Appearance Models (AAMs) [14], [10] and Constrained Local Models (CLMs) [39] are widely used techinques for keypoint tracking. Both of these methods rely on so called point distribution models (PDMs) [10] which are constructed using principal component analysis (PCA).

AAMs and CLMs usually suffer from a degree of identity-dependence. That is, a model trained on a sample of subjects does not necessarily perform well on an unseen subject. AAMs in particular suffer greatly from this effect, performing much better when samples of an individual are used for both training and testing. In this paper, we investigate the problem when no additional information about a new subject is available, as is the case in many common scenarios. As such, we evaluate our method using the AAM as a test algorithm, because it performs poorly in this setting but performs well otherwise.

View-based approaches [35], and multi-stage solutions [44], [31] address this issue by using multiple subspaces for each identity. Gaussian mixture models, [20] indirectly deal with identity variation by learning clusters of training data.

Gross *et al.* described reduced fitting robustness of AAMs on unseen subjects [22], suggesting simultaneous appearance and shape fitting improve generalization. Our method approaches identity variation directly and probabilistically, learning both factors of variation simultaneously.

Some other notable recent work from Jeni et al. has extended CLMs to account for 3D shape. Their work has demonstrated extremely high performance on the CK+ expression classification and action unit detection evaluations. We compare our approach with their method in Section 3.2.

Other work has used a structured max-margin approach to keypoint placement [56] yielding state of the art results for facial keypoint placement. We compare with this approach in Section 5.

Deep learning has recently emerged as a method capable of yielding state of the art results for keypoint placement [25, 41]. Deep networks are however known to overfit when data set sizes are small, which is the case for many of the problems that we are addressing in our work here (ex. markerless facial performance capture imagery from helmet mounted cameras).

Other factorization techniques and our approach. Multi-linear and bilinear analysis of facial images [19], [48], [50] can model the interaction of different kinds of variation. However, with these models, a full image tensor is often needed, which can be difficult to obtain. Some approaches overcome this restriction by treating the required tensor labels as missing [12]. However, this leads to discarding data.

The problem we address in this paper is that of learning an expression representation which leverages identity information, then using the technique to improve a wide variety of expression and emotion related tasks. This is similar in spirit to the expression synthesis approaches used by Du and Lin [13] and

by Zhou and Lin [55], which generate identity-independent expression factors. Our approach generalizes this work as a framework for unsupervised learning of expression factors, building in particular on the identity-expression factorization technique of Prince et al. [37]. In the next section we present their probabilistic approach to this task, then go on to show how show how this formulation can be used to create an extremely wide variety of new techniques leading to improved performance for expression recognition and facial action recognition tasks. Subsequently, in Sections 4 and 5 we go on to extend the widely used formulations of AAMs and CLMs so as to use identity-expression factorized representations. In both cases we show that keypoint localization performance is improved over the traditional techniques which use principal component analysis as their underlying point distribution models.

3. A Model for Disentangling Identity and Expression

Face image datasets have a rich divesity due to many contributing factors of variation such as age, illumination, identity, pose, facial expression and emotion. In this literature, we consider two main sources of variation. The primary variation is due to the identity factors and the second source of variation is due to the facial expressions as well as some degree of pose deviation.

A graphical model of this approach is shown in Figure 1, where each face image $\mathbf{x}_{ij} \in \mathbf{X}$ is generated by w_i, representing the identity i, and v_{ij}, representing the j^{th} expression of identity i. The set of observed data is $\mathbf{X} = \{\mathbf{x}_{ij}\}_{i=1,\ldots,I\ ;\ j=1,\ldots,J_i}$ which contains all images of I identities, where for each each identity i, a total of J_i expressions exist. The observed data can be pixels, vectors containing keypoint locations or a concatenation of both. We use the notation $\mathbf{x}_{ij}$ to denote the

j^{th} image of the i^{th} identity in the dataset.

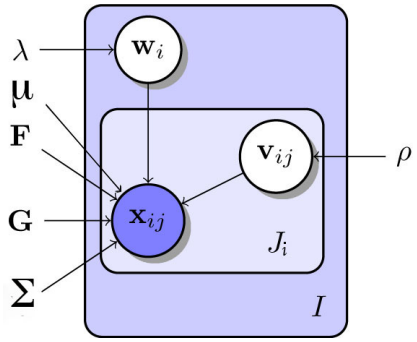

Figure 1: Graphical model of facial data generation. $\mathbf{x}_{ij}$ is generated from $p(\mathbf{x}_{ij} \mid \mathbf{v}_{ij}, \mathbf{w}_i)$, after sampling $\mathbf{w}_i$ from an identity and $\mathbf{v}_{ij}$ from an expression-related distribution respectively.

Each $\mathbf{x}_{ij}$ is generated by sampling $\mathbf{w}_i$ and $\mathbf{v}_{ij}$ from Gaussian distributions corresponding to identity i and its j^{th} expression distributions $p(\mathbf{w}_i)$ and $p(\mathbf{v}_{ij})$, and then combining these by sampling an image $\mathbf{x}_{ij}$ according to $p(\mathbf{x}_{ij} \mid \mathbf{w}_i, \mathbf{v}_{ij})$. We use zero-mean independent Gaussian distributions for $p(\mathbf{w}_i)$ and $p(\mathbf{v}_{ij})$.

$$p(\mathbf{w}_i) = \mathcal{N}(\mathbf{w}_i;\ \mathbf{0},\ \lambda\mathbf{I}), \tag{1}$$

$$p(\mathbf{v}_{ij}) = \mathcal{N}(\mathbf{v}_{ij};\ \mathbf{0},\ \rho\mathbf{I}). \tag{2}$$

The observation $\mathbf{x}_{ij}$ is then sampled from a multivariate Gaussian conditional distribution parameterized by the mean $\boldsymbol{\mu}$, matrices $\mathbf{F},\ \mathbf{G}$ and diagonal covariance $\boldsymbol{\Sigma}$, such that

$$p(\mathbf{x}_{ij} \mid \mathbf{w}_i, \mathbf{v}_{ij}) = \mathcal{N}(\mathbf{x}_{ij};\ \boldsymbol{\mu} + \mathbf{F}\mathbf{w}_i + \mathbf{G}\mathbf{v}_{ij},\ \boldsymbol{\Sigma}). \tag{3}$$

This corresponds to a conditional distribution given the variable $\mathbf{w}_i$, which is identical for all images of a unique identity, and the variable $\mathbf{v}_{ij}$, which varies across

different expressions of a particular identity. The loading matrices $\mathbf{F}$ and $\mathbf{G}$ correspond respectively to identity and expression and are shared across all observations. The joint probability can be written as

$$p(\mathbf{X}, \mathbf{W}, \mathbf{V} \mid \mathbf{F}, \mathbf{G}, \lambda, \rho, \boldsymbol{\mu}) = \prod_{i}^{I} \mathcal{N}(\mathbf{w}_i;\ \mathbf{0},\ \lambda\mathbf{I}) \prod_{j}^{J_i} \mathcal{N}(\mathbf{x}_{ij};\ \boldsymbol{\mu} + \mathbf{F}\mathbf{w}_i + \mathbf{G}\mathbf{v}_{ij},\ \boldsymbol{\Sigma}) \mathcal{N}(\mathbf{v_{ij}};\ \mathbf{0},\ \rho\mathbf{I}). \tag{4}$$

where $\mathbf{X}$ is the entire observable dataset and $\mathbf{W} = \{w_1, w_2, \ldots, w_I\}$ is the set of all latent identity representations and $\mathbf{V} = \{v_i j\}_{i=1,\ldots,I\ ;\ j=1,\ldots,J_i}$ is the set of all latent expression representations for all identities.

3.1. Learning

The Expectation Maximization (EM) approach is used to learn the parameters of the model $\boldsymbol{\theta} = \{\mathbf{F}, \mathbf{G}, \boldsymbol{\Sigma}, \boldsymbol{\mu}, \lambda, \rho\}$. The goal is to maximize the joint distribution by alternatively maximizing and taking expectations

$$\max_{\boldsymbol{\theta}} \mathbb{E}[\log p(\mathbf{X}, \mathbf{W}, \mathbf{V} \mid \boldsymbol{\theta})], \tag{5}$$

where the expectation is taken with respect to the posterior distribution

$$p(\mathbf{W}, \mathbf{V} \mid \mathbf{X}, \boldsymbol{\theta}^{\text{old}}) \tag{6}$$

and then the parameters in $\boldsymbol{\theta}$ are updated by maximizing the log of the joint distribution $p(\mathbf{W}, \mathbf{V}, \mathbf{X})$. Since the priors in Eqs. (1) and (2) and the likelihood distribution in Eq. (3) are all Gaussian, the resulting joint distribution is also Gaussian.

For a given identity i, the set of observed variables $\{\mathbf{x}_{ij}\}_{j=1,\ldots,J_i}$, can be written as a single feature vector $\mathbf{x}_i = (\mathbf{x}_{i1}^T, \mathbf{x}_{i2}^T, \ldots, \mathbf{x}_{iJ_i}^T)^T$. Similarly the factors can be combined into a single loading matrix

$$\mathbf{A}_i = \begin{pmatrix} \mathbf{F} & \mathbf{G} & 0 & ... & 0 \\ \mathbf{F} & 0 & \mathbf{G} & ... & 0 \\ \ldots & \ldots & \ldots & \ldots & \ldots \\ \mathbf{F} & 0 & 0 & ... & \mathbf{G} \end{pmatrix}, \tag{7}$$

where the number of rows and columns in $\mathbf{A}_i$ are J_i and $J_i + 1$. The latent space representation corresponding to identity i becomes

$$\mathbf{d}_i = (\mathbf{w}_i^T, \mathbf{v}_{i1}^T, \mathbf{v}_{i2}^T, \ldots, \mathbf{v}_{iJ_i}^T)^T. \tag{8}$$

Having $\mathbf{w}_i \in R^K$ and $\mathbf{v}_{ij} \in R^L$, $\mathbf{d}_i$ becomes a vector of dimensionality $K + L \times J_i$. Since $\mathbf{d}_i$ is composed of two sets of vectors with zero-mean Gaussian distributions, it is also distributed as a zero mean Gaussian:

$$p(\mathbf{d}_i) = \mathcal{N}(\mathbf{d}_i;\ \mathbf{0},\ \mathbf{\Phi_i}), \tag{9}$$

$$\mathbf{\Phi}_i = diag\left(\lambda_1,\ \ldots,\ \lambda_K,\ \rho_{1_1},\ \ldots,\ \rho_{1_L},\ \ldots,\ \rho_{Ji_1},\ \ldots,\ \rho_{Ji_L}\right), \tag{10}$$

where the term $\mathbf{\Phi}_i$ is a diagonal covariance matrix, in which the first K elements of the diagonal are extracted from the diagonal elements of the covariance matrix $\lambda\mathbf{I}$. The diagonal of $\mathbf{\Phi}_i$ is then composed of J_i blocks of L elements that are extracted from the diagonal of covariance matrix $\rho\mathbf{I}$ repeated J_i times. Intuitively the first set of K elements represent identity variations, while each of the

J_i blocks represent L factors encoding expression variations within each image. Given this construction, the probability for identity i can be rewritten as a Gaussian $p(\mathbf{x}_i \mid \mathbf{d}_i) = \mathcal{N}(\mathbf{x}_i;\ \boldsymbol{m}_i + \mathbf{A}_i\mathbf{d}_i,\ \boldsymbol{\Psi}_i)$, where $\boldsymbol{\Psi}_i$ is constructed as a diagonal matrix by concatenating J_i times the diagonal of Σ. The term $\boldsymbol{m}_i$ is J_i blocks of $\boldsymbol{\mu}$ being concatenated.

The posterior probability of $\mathbf{d}_i$ is also Gaussian with moments

$$\mathbb{E}[\mathbf{d}_i] = (\boldsymbol{\Phi}_i^{-1} + \mathbf{A}_i^T {\Psi_i}^{-1}\mathbf{A}_i)^{-1}\mathbf{A}_i^T {\Psi_i}^{-1}(\mathbf{x}_i - \boldsymbol{m}_i), \tag{11}$$

$$\mathbb{E}[\mathbf{d}_i\mathbf{d}_i^T] = (\boldsymbol{\Phi}_i^{-1} + \mathbf{A}_i^T {\Psi_i}^{-1}\mathbf{A}_i)^{-1} - \mathbb{E}[\mathbf{d}_i]\mathbb{E}[\mathbf{d}_i^T]. \tag{12}$$

In the E-step, the mean and covariance matrices of the posterior distribution $p(\mathbf{d}_i \mid \mathbf{x}_i)$ is taken, which is measured over all of the expressions corresponding to the same identity. This representation style assures that all of the expressions of the same person have the same representation for identity. Note that the expectations in Eqs. (11) and (12) should be taken separately for all of the identities in the training set. In the M-step, however, we can disentangle the $\mathbf{d}_i$ into a set of $\mathbf{d}_{ij}$, in which each $\mathbf{d}_{ij}$ contains one sample of identity and one sample of expression. This is due to the fact that the parameters of the model should be updated with respect to all of the data in the training set and having done the E-step, we have already got the same identity representation for all expression observations from the same person. Therefore, disentangling $\mathbf{d}_i$ into a set of $\mathbf{d}_{ij}$ can reserve that information while at the same time it encodes a simpler latent representation for updating the parameters of the joint distribution. We define $\mathbf{d}_{ij} = \begin{bmatrix} \mathbf{w}_i \\ \mathbf{v}_{ij} \end{bmatrix}$ and

$\mathbf{B} = [\ \mathbf{F} \quad \mathbf{G}\ ]$, such that d_{ij} is distributed as

$$p(\mathbf{d}_{ij}) = \mathcal{N}(\mathbf{d}_{ij}; \mathbf{0}, \mathbf{\Phi}), \tag{13}$$

$$\mathbf{\Phi} = diag\left(\lambda_1, \ldots, \lambda_K, \rho_1, \ldots, \rho_L\right). \tag{14}$$

Given this notation, the conditional distribution can now be written as

$$p(\mathbf{x}_{ij}|\mathbf{d}_{ij}) = \mathcal{N}(\mathbf{x}_{ij}; \boldsymbol{\mu} + \mathbf{B}\mathbf{d}_{ij}, \mathbf{\Sigma}). \tag{15}$$

We can further simplify this notation by setting $\tilde{\mathbf{d}_{ij}} = \begin{bmatrix} \mathbf{d}_{ij} \\ 1 \end{bmatrix}$ and $\tilde{\mathbf{B}} = [\ \mathbf{B} \quad \boldsymbol{\mu}\ ]$, which in turn gives the following conditional and prior distributions

$$p(\mathbf{x}_{ij}|\tilde{\mathbf{d}_{ij}}) = \mathcal{N}(\mathbf{x}_{ij}; \tilde{\mathbf{B}}\tilde{\mathbf{d}_{ij}}, \mathbf{\Sigma}), \tag{16}$$

$$p(\tilde{\mathbf{d}_{ij}}) = \mathcal{N}(\tilde{\mathbf{d}_{ij}}; \mathbf{0}, \tilde{\mathbf{\Phi}}) \tag{17}$$

with $\tilde{\mathbf{\Phi}}$ being equal to

$$\tilde{\mathbf{\Phi}} = \begin{bmatrix} \mathbf{\Phi} & \mathbf{0_{K+L,1}} \\ \mathbf{0_{1,K+L}} & \mathbf{0_{1,1}} \end{bmatrix}, \tag{18}$$

where two zero vectors with row or column size of $K + L$ are concatenated with $\mathbf{\Phi}$ and another zero element to build a square matrix of size $K + L + 1$. The joint

distribution then gets equal to

$$p(\mathbf{X}, \boldsymbol{\mathcal{D}}|\boldsymbol{\Sigma}, \tilde{\mathbf{B}}, \tilde{\boldsymbol{\Phi}}) = \prod_{\mathbf{ij}} \mathcal{N}(\mathbf{x_{ij}}; \tilde{\mathbf{B}}\tilde{\mathbf{d_{ij}}}, \boldsymbol{\Sigma}) \mathcal{N}(\tilde{\mathbf{d_{ij}}}; \mathbf{0}, \tilde{\boldsymbol{\Phi}}). \tag{19}$$

with $\mathbf{X}$ representing the whole set of observed data and $\boldsymbol{\mathcal{D}}$ representing the entire set of latent variables. This simplification is preferred since the variables $\mathbf{F}, \mathbf{G}, \boldsymbol{\mu}$ are mutually dependent. Updating each one requires the other two. In the new notation, all of them can be updated simultaneously by updating $\tilde{\mathbf{B}}$. Maximizing with respect to $\tilde{\mathbf{B}}, \boldsymbol{\Sigma}, \tilde{\boldsymbol{\Phi}}$, give the following updates

$$\tilde{\boldsymbol{\Phi}} = \frac{1}{N} \sum_{ij} \text{diag} \left\{ \mathbb{E} \left[\tilde{\mathbf{d}}_{ij} \tilde{\mathbf{d}}_{ij}^T \right] \right\}, \tag{20}$$

$$\boldsymbol{\Sigma} = \frac{1}{N} \sum_{ij} \text{diag} \left\{ \tilde{\mathbf{B}} \mathbb{E} \left[\tilde{\mathbf{d}}_{ij} \tilde{\mathbf{d}}_{ij}^T \right] \tilde{\mathbf{B}}^T + \mathbf{x}_{ij} \mathbf{x}_{ij}^T - 2 \mathbf{x}_{ij} \mathbb{E} \left[\tilde{\mathbf{d}}_{ij}^T \right] \tilde{\mathbf{B}}^T \right\}, \tag{21}$$

$$\tilde{\mathbf{B}} = \sum_{ij} \left\{ \mathbf{x}_{ij} \mathbb{E} \left[\tilde{\mathbf{d}}_{ij}^T \right] \right\} \left\{ \sum_{ij} \mathbb{E} \left[\tilde{\mathbf{d}}_{ij} \tilde{\mathbf{d}}_{ij}^T \right] \right\}^{-1}. \tag{22}$$

Note that the parameters are updated with respect to all expressions of all identities. As for inference at test time, the procedure for determining the optimal $\mathbf{d}_{ij}$ vector is straight-forward, using the following posterior distribution:

$$p(\mathbf{d}_{ij}|\mathbf{x}_{ij}) = \mathcal{N}(\mathbf{d}_{ij}; (\boldsymbol{\Phi}^{-1} + \mathbf{B}^T \boldsymbol{\Sigma}^{-1} \mathbf{B})^{-1} \mathbf{B}^T \boldsymbol{\Sigma}^{-1} (\mathbf{x}_{ij} - \boldsymbol{\mu}), (\boldsymbol{\Phi}^{-1} + \mathbf{B}^T \boldsymbol{\Sigma}^{-1} \mathbf{B})^{-1}). \tag{23}$$

3.2. Facial Expression Analysis, Animation Control Methods and Experiments

In the following sub-sections, we build upon the model above to create techniques for emotion recognition and performance-driven animation and evaluate

the methods using multiple datasets. We first present emotion recognition methods and results, including the task of facial action unit classification. Then, we move on to performance-driven animation techniques and experiments.

A summary of these experiments is shown in Table 1.

3.2.1. Emotion Recognition with JAFFE

Figure 2: The JAFFE dataset contains 213 labeled examples for 10 subjects. Images for a single test subject, left out from training, shown here with predicted labels from our method. The data is shown in two sets of rows. The first row in each set is the original input data, the second a rendering of the mesh with corresponding bone positions predicted by the model.

We first run a set of experiments on the constrained JAFFE dataset [33]. The JAFFE dataset contains ten identities with varying expression across seven emotions containing 213 images in total. The frontal facial images are roughly aligned but we use a funneling algorithm [26] after face detection to correct small pose variation.

We then use an ellipse, manually specified, to mask background variation and reduce dimensionality, and then divide the face into three rectangular regions roughly corresponding to the mouth, eyes and ears. However, unlike previous works, such as [55], we do not manually locate facial landmarks. Then we run the

algorithm for each region individually, learning a composite space in which facial expressions are treated as a linear combinations. The expression space weights for each region are learned separately and then concatenated to form a single input vector consisting of the weights $\mathbf{v}_{ij}$.

We then predict seven emotional states – anger, disgust, fear, happiness, neutral, sadness and surprise – using the learned representation $\mathbf{v}_{ij}$. We train a SVM RBF classifier for each of seven emotional states on the in-sample subject images, using a single image from each of the emotions of the left-out subject as a validation set to learn the slack and kernel bandwidth parameters. We treat the classification at test time as a one-against-all prediction. We compare our results against PCA performed on the same input data as our model, using 30 and 100 dimensions. The final rates are shown in Table 2 as the average accuracy across identities.

The overall prediction accuracy rate for the emotion prediction for an unseen identity is 72.17% using a 30 dimensional expression space. A recent result on the JAFFE set by Cheng *et al.* [8], in experiments with left-out subjects, obtained a mean accuracy of 55.24% using Gaussian process classification. Because their pre-processing is not identical and they do not report standard errors, we do not include it in Table 2.

3.2.2. Emotion Recognition and Facial Action Unit Detection with CK+

Certainly, the JAFFE data exhibits far less variation than is usually present in real-world data. In this section, we present more detailed experiments using the Extended Cohn Kanade (CK+) database [32] for emotion recognition and facial action unit tasks. The CK+ dataset consists of 593 image sequences from 123 subjects ranging in age from 18 to 50, 69% of whom are female and 13% of whom

are black. The images are frontal images of posed subjects taken from video sequences. Each sequence contains a subject posing a single facial expression starting from a neutral position. The sequence consists of sampled frames from videos in which the final posed position is labeled with FACS action units. In addition, emotion labels, consisting of the expressions "Anger," "Disgust," "Fear," "Happiness," "Sadness," "Surprise" and "Contempt", are provided for 327 of the 593 sequences.

Facial Action Unit Detection. The CK+ database contains 593 labeled sequences, however only the final image as well as the initial neutral image may be used for traditional classification. Lucey *et al.* [32] describe their baseline approach based on linear classification of key-points and warped images obtained by Active Appearance Models. The resulting landmark data is also provided in this dataset. To compare our method with the baseline, we recreate their approach. First, we use a Procrustes analysis using affine transformation of the landmark positions to determine a mean shape and register the point locations to the mean by estimating the least square best 2D transformation. We then run PCA on the difference between the Procrustes aligned points and the original points to determine the similarity components. We add these to the AAM shape parameters to model rigid motion. Finally, a piece-wise affine warp is applied to normalize the shape of each facial image to the base shape recovered from the Procrustes analysis to obtain warped images.

Using this approach, we generate two feature sets. The first is point locations after Procrustes analysis. The second set contains shape-normalized images, which are converted to vectors. We use leave-one-subject-out cross-validation using linear SVM's for each of the 17 AU's on the vectors of landmark locations and

vector of shape-normalized images, and record the AUC score for each left-out subject and AU pair. An estimate for the AUC error is calculated as well, defined as $\sqrt{\frac{A(1-A)}{\min(N_p,N_n)}}$, where A denotes the AUC score and N_p, N_n are the number of positive and negative examples respectively. To combine feature sets, Lucey *et al.* run logistic regression on the scores of SVM's built from the two feature-sets independently. We also recreate this step.

For our identity-normalization experiments, we use the point-locations and shape-normalized images as inputs to our factor analysis. In this case, we use 100 and 30 dimensions for the identity and expression parameter vectors respectively. In order to avoid overfitting, we increase the training data size for this unsupervised step from 1186 to 2588 by ensuring that each identity has at least 20 images, by sampling uniformly from intermediate frames. We do not use any testing subject images during training our models, including during the factor analysis in order to maintain fairness. The results are shown in Table 3, which shows an increase in average AUC for each of the feature-sets individually. In a majority of AU classes, the AUC score increased, while there are no large reductions in AUC after identity normalization. For both the SPTS and CAPP scores, identity normalization showed significant increases in performance for AU-7, Lid-Tightener and AU-15, Lip Corner Depresser. AU-7 is associated with an eye-narrowing or squinting action, and AU-15. We show examples from this dataset of these two AU's to show that resulting normalization leads to representations that can be seen as roughly identical.

In both cases, the expression is well reproduced in the identity-normalization procedure. The increase in performance from using both feature-sets is apparent again, which is greater than .86 AUC for every AU and well over 90% for most

of them. Our comparison shows that the simple step of identity normalization improves performance on average. It appears to work especially well for those AU's that are easy, in some sense, to capture between identities.

We also place our work in further context with the method recently proposed by Jeni et al. [28] for shape based action unit recognition. They give an average AUC result of 89.13%. Our shape, texture and combined average AU results were 92.8%, 94.4% and 96.1%, respectively.

Emotion Detection. The baseline approach can also be used for emotion recognition. Again, we recreate their approach and use the identity-normalization step as an unsupervised learning procedure. Because of this, we can use more training examples than those that are labeled, making this procedure semi-supervised. Again, we use the 2588 images used in the AU experiments. Each of the 327 labeled examples along with the neutral examples are then projected into the space learned in the previous step and used as input for training the linear SVM's. In this case, each SVM learns a one-vs-all binary classifier for the emotion of interest, using the rest as negative examples. The multi-class decision is made using the maximum score. In order to recreate the experiments in [32], the neutral examples are left out of the testing, resulting a forced-choice between the seven emotions. Again, the SVM's and logistic regressors are learned using leave-one-subject-out, that is, over a total of 123 trials.

However, after combining scores from the SVM's using a logistic regression technique, the improvements are quite impressive. As in [32], results after the calibration of SVM scores indicate that the point data and image data capture different kinds of information. Combining these features yields impressive gains in performance. "Anger," especially improves using our method. In particular, only

"Happiness" recognition is reduced using our approach, but the reduction in hit rate is modest. Overall, this method gave a 95.21% average accuracy compared to the result of 83.27% reported in [32]. Comparison of confusion matrices of emotion detection for the combined landmark (SPTS) and shape-normalised image (CAPP) features before and after identity normalisation is shown in Table 4.

More recently Jeni et al. [28] have proposed a 3D shape model and explored two normalization methods for the CK+ expression recognition evaluation. They report that their action unit 0 or AU0 normalization procedure yields an average class accuracy of 96%, while their personal mean shape procedure yielded an average class accuracy of 94.8%. Our average accuracy above is given as the per example average whereas these results were computed on a per-class basis. Our average per-class accuracy was 93.1%.

3.2.3. Animation Control Experiments with JAFFE

We again begin using the JAFFE dataset. Each image, after face detection and alignment, is labeled using a common 3D face mesh fitted with 27 bones created by an artist, as can be seen in Figure 2. Each bone is then positioned by an artist to a maximal and minimal position along a fixed path. The position along the path is specified by a real value in $[-1, 1]$ as a fraction of the distance between the midpoint and the extreme values.

We use MSE for evaluation purposes on the animation experiment, which corresponds to bone parameter recovery. In this case, we use linear regression as the predictive algorithm. The MSE reported is the average error in bone position for all test images. For each trial, we leave one identity out from the trial, and compute both the MSE error and a prediction for the facial action. We compare our approach to PCA, using 30 and 100 dimensions. The results are summarized

in Table 2. For the experiment labeled "none," the experiment denotes no unsupervised pre-processing step – the input is the raw image data. As it is difficult to gauge the quality of the predictions from MSE alone, the predictions for a test subject using our method are illustrated in Figure 2, showing that the method does recover the facial actions produced by the subject quite well. The small standard error of the MSE indicates that, in general, the procedure is capable of predicting facial actions across all unseen subjects.

3.2.4. Animation Control Experiments with Markerless Motion Capture

We now return to the challenging real-world problem of high quality facial animation control using video obtained from helmet cameras. This technique was used in the well known film Avatar and this data comes from Ubisoft, the company responsible for the brands Assassin's Creed and FarCry among others. FarCry 3 uses these technologies extensively. We obtained 16 videos of motion capture data without marker data. These videos are produced using infra-red cameras. Examples are shown in Figure 3. As is common in motion capture data, faces are captured using a helmet-mounted camera which reduces pose variation. The camera position is relatively fixed. However, this data presents new challenges due to variation stemming from the varied appearance of the actors. To compensate for appearance variation due to facial structure, an Active Appearance Model is applied to the data as shown in Figure 3, using 66 fiducial points. The resulting points derived from the model are used to warp each video frame to a common coordinate structure, removing pixels from outside the convex hull. These were used as inputs to our model, for a dataset totaling 3122 frames. Example outputs can also be seen in Figure 3b.

For our evaluations, we were given professional results of a bone-based model

(a) Original data

(b) Results of AAM model

(c) Piecewise warped data

(d) Identity normalization is applied

Figure 3: Motion capture training data using a helmet IR camera. Example pipeline for a test video sequence.

used as motion capture output. The bone positions are given values in $[-1, 1]$ representing the relative distance from the neutral position for each bone along predefined space. Each experiment is run with one subject left out, and the results shown are the average MSE and standard error. Results of the experiments are shown in the right-most column of Table 2. We first show results using no preprocessing except for face detection using the Viola-Jones detector, cropped to a size 1.5 time the size of the detector to account for the distortion evident in the images. The results shown in Table 2 are per bone, and clearly some data processing must be used (None* refers to the images after shape normalization). Best results were not obtained by applying PCA to the original data. The remaining experiments

used shape-normalised data as inputs. As can be seen PCA improves the results, but the best results are achieved using our model. The classifier used was a simple linear regression, a more powerful discriminative classifier may be more effective.

4. Identity-Expression Active Appearance Models (IE-AAMs)

For many of the experiments in the preceding section a key step is in the application of key-point detection algorithms, specifically AAMs. The AAM has been shown to work very well for subjects on which it has been trained while generalization to unseen identities has proven more difficult, due to the increased complexity of the shape model [22]. Since labeling data for new identities can be labor intensive we would like to determine if AAM performance can be increased for an unseen identity using an identity-expression factorization extension to the traditional AAM formulation.

The traditional formulation on an AAM contains a PCA model of the joint point and appearance data. This lends itself to a linear Gaussian interpretation which leads us to investigate whether key-point localization itself is improved using identity normalization. In this section we develop a novel identity-expression factorized AAM and investigate the utility of identity normalization for key point detection.

4.1. Inverse Compositional AAMs

The AAM fitting procedure we adapt is the inverse-compositional method as described by Matthews and Baker in [34]. We briefly review this method here. The modified objective in AAM fitting is to minimize the error between a template image, T and a set of warped images $I_j, j = 1, 2, ..., J$, which generally belong to the same identity. The modified objective is to minimize the squared error between

the set of image over all image locations $\mathbf{x} = (x, y)$, by estimating parameters $\mathbf{p}_j$ of a warp function W, or

$$\min_{\mathbf{p}_j} \sum_j \sum_{\mathbf{x}} [T - I_j(W(\mathbf{x}; \mathbf{p}_j))]^2. \tag{24}$$

We change notation here to make the following more clear. The warp function, W, can be interpreted as a map which determines a new location for any point $\mathbf{x}$, *i.e.* , $\mathbf{x}' = W(\mathbf{x}, \mathbf{p}_j)$. For AAM models, this is typically a piecewise affine transform based on the triangulation of a shape $\mathbf{s}$, which are represented by (x, y) point locations.

We shall quickly review how this optimization is performed with the classical AAM inverse compositional mapping approach, then replace the standard approach with our identity and expression factorization method and see the impact on the optimization procedure. In the standard AAM, $\mathbf{p}_j$ is the vector of weights for a set of eigen-shapes. The shape $\mathbf{s}$ is parameterized by a PCA model such that each $\mathbf{s}_j$ is explained by the mean shape $\mathbf{s}_0$ and a weighted combination of eigen-vectors $\mathbf{P}_\mathbf{s}$ with weight parameters $\mathbf{p}_j$, such that $\mathbf{s}_j = \mathbf{s}_0 + \mathbf{P}_s\mathbf{p}_j$.

The inverse compositional approach is to optimize this objective by iteratively building up a series of warps to recover an optimal warp $W(\mathbf{x}, \mathbf{p}_j)$. Somewhat confusingly, both the template, T and the image I are warped. T is always warped from its initial shape $\mathbf{s}_0$. I, however, is warped to the current estimate $I(W(\mathbf{x}, \mathbf{p}_j))$. Once the optimal warp at the iteration is determined, $\Delta\mathbf{p}_j$, the current warp parameters $\mathbf{p}_j$ and $\Delta\mathbf{p}_j$ are "composed" in order to update the current warp parameters. When $\Delta\mathbf{p}_j$ is close to zeros, or if the warp is not changing much, the algorithm has converged. The current warp is formed by inverting the

parameters of $\Delta\mathbf{p}_j$ (by negating) and applying the warp defined by these updated parameters to a warp composed of all previous warps. The modified objective at each iteration is then

$$\min_{\mathbf{p}_j} \sum_j \sum_{\mathbf{x}} [T(W(\mathbf{x}; \Delta\mathbf{p}_j)) - I_j(W(\mathbf{x}; \mathbf{p}_j))]^2. \tag{25}$$

Using a first order Taylor expansion, we have

$$\min_{\mathbf{p}_j} \sum_j \sum_x [T(W(\mathbf{x}; 0)) + \nabla T \frac{\partial W}{\partial \mathbf{p}} \Delta\mathbf{p}_j - I_j(W(\mathbf{x}; \mathbf{p}_j))]^2. \tag{26}$$

So that $\Delta\mathbf{p}_j$ can be given as

$$\Delta\mathbf{p}_j = \mathbf{H}^{-1} \sum_{\mathbf{x}} \left[\nabla T \frac{\partial W}{\partial \mathbf{p}}\right] [I_j(W(\mathbf{x}; \mathbf{p}_j)) - T(\mathbf{x})], \tag{27}$$

for an individual $\mathbf{p}_j$, where $\mathbf{H}$, the Hessian, can be written as

$$\mathbf{H} = \sum_{\mathbf{x}} \left[\nabla T \frac{\partial W}{\partial \mathbf{p}}\right]^T \left[\nabla T \frac{\partial W}{\partial \mathbf{p}}\right]. \tag{28}$$

The algorithm consists of pre-computing ∇T, $\frac{\partial W}{\partial \mathbf{v}_j}$, $\mathbf{H}_v, \mathbf{H}_w$ and inverses at the mean shape for efficiency. These are used along with Eq. (27) to compute the change at $(\mathbf{p}_j = 0)$ required to minimize the error between the template and the current warped target. To adapt this algorithm for our purposes requires few changes.

4.2. A Novel Identity-Expression AAM (IE-AAM) Formulation

In the AAM inverse compositional method, one must determine $W(\mathbf{x}, \mathbf{p}^k) = W(\mathbf{x}, \mathbf{p}^{k-1}) \circ W(\mathbf{x}; \Delta\mathbf{p})^{-1} \approxeq W(\mathbf{x}, \mathbf{p}^{k-1}) \circ W(\mathbf{x}; -\Delta\mathbf{p})$ at each iteration k. To compute the parameters of $W(\mathbf{x}, \mathbf{p}^k)$ one then computes the corresponding changes to the current mesh vertex locations $\Delta\mathbf{s} = (\Delta x_1, \Delta y_1, \ldots, \Delta x_v, \Delta y_v)^T$, which are computed by composing the current warp.

The new parameters at time k are then given by $\mathbf{p}_j^k = \mathbf{P}_s^{-1}(\hat{\mathbf{s}}_j)$, where $\hat{\mathbf{s}}_j = \mathbf{s}_j^{k-1} + \Delta\mathbf{s}_j - \mathbf{s}_0$. Now, in the case of our identity and expression decomposition model, we replace the PCA model with a factorized model with the form $\mathbf{s}_j = \mathbf{s}_0 + \mathbf{F}\mathbf{w} + \mathbf{G}\mathbf{v}_j$. The parameters are the vectors $\mathbf{w}$ and $\mathbf{v}_j$ belonging to the identity and the j^{th} expression of that identity. Note that compared to the notation used in Section 3.1, we have omitted the index i since we are dealing with the same identity. In our identity and appearance decomposition approach, we now have an objective of the form,

$$\min_{\mathbf{w},\mathbf{v}_j} \sum_j \sum_{\mathbf{x}} [T(\mathbf{W}(\mathbf{x}; \Delta\mathbf{w}, \Delta\mathbf{v}_j)) - I_j(\mathbf{W}(\mathbf{x}; \mathbf{w}, \mathbf{v}_j))]^2. \tag{29}$$

Defining $\mathbf{z}_j = [\mathbf{w}^T, \mathbf{v}_j^T]^T$, Taylor series expansion yields

$$\min_{\mathbf{w},\mathbf{v}_j} \sum_j \sum_{\mathbf{x}} [T(\mathbf{W}(\mathbf{x}; 0, 0)) + \nabla T \frac{\partial \mathbf{W}}{\partial \mathbf{z}_j}(\Delta\mathbf{z}_j) - I_j(\mathbf{W}(\mathbf{x}; \mathbf{z}_j))]^2, \tag{30}$$

where $\Delta\mathbf{z}_j = [\Delta\mathbf{w}^T, \Delta\mathbf{v}_j^T]^T$. As identity and expression are assumed to be independent, this can be computed separately,

$$\Delta\mathbf{v}_j = \mathbf{H}_{\mathbf{v}}^{-1} \sum_{\mathbf{x}} \left[\nabla T \frac{\partial \mathbf{W}}{\partial \mathbf{v}} \right] [I_j(\mathbf{W}(\mathbf{x}; \mathbf{w}, \mathbf{v}_j)) - T(\mathbf{x})],$$

$$\Delta \mathbf{w} = \mathbf{H}_{\mathbf{w}}^{-1} \sum_{\mathbf{x}} \left[\nabla T \frac{\partial \mathbf{W}}{\partial \mathbf{w}} \right] [\frac{1}{J} \sum_{j} (I_j(\mathbf{W}(\mathbf{x}; \mathbf{w}, \mathbf{v}_j)) - T(\mathbf{x}))].$$

We then compute the warp composition as in [34],

$$W(\mathbf{x}, \mathbf{w}^k, \mathbf{v}_j^k) = W(\mathbf{x}, \mathbf{w}^{k-1}, \mathbf{v}_j^{k-1}) \circ W(\mathbf{x}; -\Delta \mathbf{w}, -\Delta \mathbf{v}_j).$$

However, computing the parameters $\mathbf{w}^k$ and $\mathbf{v}_j^k$ from $\Delta \mathbf{s}_j$ is no longer a simple matrix multiplication, as was the case when we updated $\mathbf{p}_j$ above, because the loading matrices of the identity expression model $\mathbf{F}$ and $\mathbf{G}$ are not orthonormal and we use the sequence of images. Therefore, we use the expectation of Eq. (23). That is, after computing Δs_j for each warp, let

$$\mathbf{d} \triangleq (\mathbf{w}^T, \mathbf{v}_1^T, \mathbf{v}_2^T, ..., \mathbf{v}_J^T)^T, \tag{31}$$

$$\hat{\mathbf{s}} \triangleq ((\mathbf{s}_1^{k-1} + \Delta \mathbf{s}_1 - \mathbf{s}_0)^T, ..., (\mathbf{s}_J^{k-1} + \Delta \mathbf{s}_J - \mathbf{s}_0)^T)^T. \tag{32}$$

Then, with $\mathbf{\Psi}$ constructed as a diagonal matrix with $\mathbf{\Sigma}$ along the diagonal repeated J times, $\mathbf{A}_i$ as defined in Eq. (7), and $\mathbf{\Phi}$ constructed as a diagonal matrix, by merging the diagonal of $\lambda \mathbf{I}$ with J blocks of diagonal of $\rho \mathbf{I}$ we can get

$$\mathbb{E}[\mathbf{d}_i] = (\mathbf{\Phi}^{-1} + \mathbf{A}^T \Psi^{-1} \mathbf{A})^{-1} \mathbf{A}^T \Psi^{-1} (\hat{\mathbf{s}}^k) \tag{33}$$

with $\mathbf{\Phi}$ and $\mathbf{A}$ being equal to

$$\mathbf{\Phi} = diag\left(\lambda_1, \ \ldots, \ \lambda_K, \ \rho_{1_1}, \ \ldots, \ \rho_{1_L}, \ \ldots, \ \rho_{J_1}, \ \ldots, \ \rho_{J_L}\right), \tag{34}$$

$$\mathbf{A} = \begin{pmatrix} \mathbf{F} & \mathbf{G} & 0 & ... & 0 \\ \mathbf{F} & 0 & \mathbf{G} & ... & 0 \\ \dots & \dots & \dots & \dots & \dots \\ \mathbf{F} & 0 & 0 & ... & \mathbf{G} \end{pmatrix} \tag{35}$$

from which we can then compute the current warp $W(\mathbf{x}, \mathbf{w}^k, \mathbf{v}_j^k)$.

Texture variation is handled by alternating between training the model for texture variation and estimating the texture independently. Alternating between the two parameter optimization yielded the best results, although more efficient methods can be applied.

4.3. Keypoint Localization Experiments with CK+ and IE-AAMs

Since the CK+ dataset is also supplied with facial landmark annotations, we experimented using the Active Appearance Model extension described in previous sections. We used the same training set as described in the emotion recognition sections, training the factorized model on all but a single identity, and testing on those belonging to a single identity. For all our experiments, we use leave-one-subject-out cross-validation of a dataset of 2588 images and point locations. On average, this resulted in 2566 training images per experiment. The test-set were comprised of all remaining images from the CK+ dataset.

We evaluate our approach by adapting an existing software package, the ICAAM software package [49] implemented in Matlab, replacing the PCA point distribution model with our factorized model, and comparing the results against the original software. Although more complex AAM software is available (for example the ICAAM package, which is not multi-resolution), our experiments are designed to illustrate both the effectiveness of the model and the simplicity of

adapting existing approaches.

In order to achieve good results using this package, however, we used some pre-processing which improved ICAAM results: we first applied the OpenCV Viola-Jones frontal face-detector to remove large-scale pose variation and then aligned the images using a Procrustes analysis. The variation after Procustes alignment was modeled using PCA with 3 components and added to the shape projection matrix in ICAAM in order to model pose. In our case, we added the pose parameters to the non-identity component matrix $\mathbf{G}$ and then ortho-normalization was applied, which follows Lucey *et al.* among others [32]. ICAAM was set to use the default 98% of shape variation, while our method used 200 components for identity and 50 for expression. Test images were initialized using the face-detector used for training.

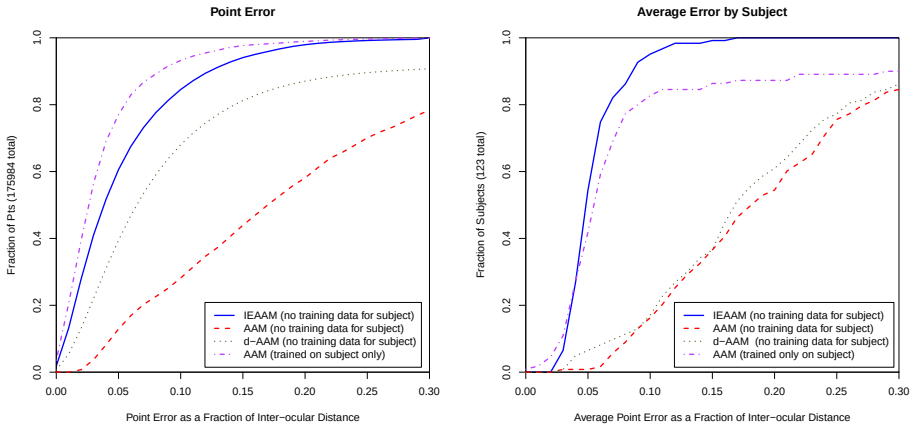
Point Error
Fraction of Pts (175984 total)
Point Error as a Fraction of Inter−ocular Distance
IEAAM (no training data for subject)
AAM (no training data for subject)
d−AAM (no training data for subject)
AAM (trained on subject only)
Average Error by Subject
Fraction of Subjects (123 total)
Average Point Error as a Fraction of Inter−ocular Distance
IEAAM (no training data for subject)
AAM (no training data for subject)
d−AAM (no training data for subject)
AAM (trained only on subject)

(a) Cumulative error distribution.

(b) Subject average error, as fraction of inter-ocular distance, cumulative distribution.

Figure 4: Point-localization experiment evaluation, described in detail in text.

The results are summarized in Figure 4. In 4a we show the total cumulative distribution of error in all tested points for four methods. This analysis is

widely used to evaluate keypoint techniques [34]. The Identity-Expression AAM, (IEAAM), out-performs a discriminative AAM included in the Demolib distribution by Saragih *et al.* [27] (d-AAM) and the ICAAM AAM on which our INM is based. In both cases, the default settings are used, with PCA dimensionality chosen to account for 95% and 98% of variation. These three methods are not trained on the left-out subject, but on all remaining identities.

To provide another point of reference concerning what AAM is capable of achieving if training data is used for the subject of interest, we show the results of a subject-specific model using the AAM implementation provided in the DeMoLib distribution. This subject-specific AAM was trained on one sample of testing subject images described previously in Section 3.2.2, and tested on all remaining subject images from the CK+ dataset. This corresponds to using the first and last images, along with a sampling from intermediary images in each sequence as training images. In CK+, this corresponds to extreme expression poses. This configuration simulates the labeling of a range of motion image sequence, which can be a good practical strategy for fitting an AAM with a small amount of labeling effort. The subject-specific AAM uses the project-out approach, [34] and the inverse-compositional method. As expected, this method performs well, with less than 22.9% of predicted points being greater than 5% of the inter-ocular distance (calculated as the two farthest points on the eyes) from the ground truth.

We also compute the average inter-ocular distance error for each subject. In 4b, we show the cumulative distribution of average subject error, which evaluates how well methods perform within each subject trial. The subject-trained AAM suffers from convergence errors, *i.e.* dramatic failures where all or many keypoints completely fall off the face, which cause the average error for each subject to

increase. In fact, the average error of all points for the subject-trained AAM is 23.3% of inter-ocular distance, while the average error of our method is 5.94%. Our method uses joint optimization on a sequence of frames adding stability to the point localization. Because of this, our method does not have these kinds of convergence errors, as shown in Figure 4b,

Other methods, such as the one of Van der Maaten and Hendriks report 4.69 pixel error using mixture modeling on the CK+ data [47]. They also report 0.23% convergence error rate at that level of accuracy. More crucially, however, test subjects in their testing regimen were not strictly prohibited from the training set. We are not aware of any AAM experiments using a large number of disparate subjects and testing on left-out subjects that report lower error. In any case, we are motivated more by the possibility of improving methods by using identity information, and believe that identity-normalization might easily be applied to any AAM-method. However, we do note that we see 0% convergence error, which we believe is due to the inference over multiple images of the same identity.

These results strongly suggest that identity is an important source of information and there are measurable benefits to modeling such information explicitly. The average error for the CK+ set using the ICAAM code with face detection initialization is 24.36, with a standard deviation of 11.10, whereas our model yields an average pixel error of 7.15, with a standard deviation of 6.57. Both the bias and the variance is minimized simply by extricating the source of variation. The change to the algorithm, as discussed in previous sections, is relatively small. The result is significantly improved.

5. Identity-Expression Constrained Local Models (IE-CLMs)

The term Constrained Local Models (CLMs) has evolved from the original work of Cristinacce and Cootes [11], which can be viewed as a particular instance of a CLM [39]. Nowadays the term CLM has come to refer to a number of methods which involve finding the landmarks of an image $\mathcal{I}$ through assigning a cost $\mathcal{Q}$ to candidate landmark positions $\{\mathbf{x_1}, ..., \mathbf{x_n}\} \in \mathbf{x}$ on the image and the parameters of the model $\mathbf{p}$. The corresponding objective function can be written as:

$$\mathcal{Q}(\mathbf{x}, \mathbf{p}) = \sum_{\mathbf{i=1}}^{\mathbf{n}} \mathcal{D}_{\mathbf{i}}(\mathbf{x_i}; \mathcal{I}) + \mathcal{R}(\mathbf{x}, \mathbf{p}), \tag{36}$$

where $\mathcal{D}_i$ encodes the image dependent suitability measure for the i^{th} landmark being located at position $\mathbf{x_i}$ in the image. The $\mathcal{R}$ term can be interpreted as a regularization term that encodes preferences for certain spatial configurations of landmark positions[1]. This set-up leads to what some refer to as a deformable model fitting problem. Further, it is common to use a linear approximation for how the shape of non-rigid objects deform and a common variant of such a modeling technique is known as a point distribution model or PDM [10]. The PDM of [10] models non-rigid shape variations linearly and composes them with a global rigid transformation such that the 2D location of the PDM's i^{th} landmark is given by:

$$\mathbf{x}_i = s\mathbf{R}(\boldsymbol{\mu}_i + \mathbf{\Phi_i q}) + \mathbf{t}, \tag{37}$$

[1] We have reversed the order of terms compared to the notation in [39] and indicated an explicit dependence on $\mathbf{x_i}$ for the underlying objective and $\mathcal{R}$.

where the PDM parameters are defined by $\mathbf{p} = \{s, \mathbf{R}, \mathbf{t}, \mathbf{q}\}$. These parameters consist of a global scale s, rotation $\mathbf{R}$, and translation $\mathbf{t}$ (forming a similarity transformation), as well as a set of parameters encoded in vector $\mathbf{q}$ which capture the global non-rigid deformations though a sub-matrix $\mathbf{\Phi_i}$ of a larger matrix $\mathbf{\Phi}$ capturing variations using basis vectors. Using a probabilistic notation, the probability over the landmark position $\mathbf{x}_i$ can be represented as a normal distribution with mean $s\mathbf{R}(\boldsymbol{\mu}_i + \mathbf{\Phi_i q}) + \mathbf{t}$ and a covariance of σ^2, such that

$$p(\mathbf{x_i}) = \mathcal{N}(s\mathbf{R}(\boldsymbol{\mu_i} + \mathbf{\Phi_i q}) + \mathbf{t}, \boldsymbol{\sigma^2}). \tag{38}$$

Consequently, the distribution over the whole set of landmarks $\mathbf{x}$ gets the form of a normal distribution which corresponds to

$$p(\mathbf{x}) = \mathcal{N}(s\mathbf{R}(\boldsymbol{\mu} + \mathbf{\Phi q}) + \mathbf{t}, \boldsymbol{\sigma^2 I}), \tag{39}$$

where $\boldsymbol{\mu}$ and $\mathbf{\Phi}$ represent respectively a concatenation of the terms $\boldsymbol{\mu}_i$ and $\mathbf{\Phi}_i$ for $i \in \{1, \ldots, n\}$. The Covariance $\sigma^2\mathbf{I}$ is a diagonal matrix having the values in its diagonal being repeated n times. Eq. (39) gives a principled probabilistic representation for the regularization term $\mathcal{R}$ in Eq. (36). It is common to place a uniform prior on the parameters of the similarity transformation, while the model for non-rigid motion is typically represented using principal component analysis.

5.1. A Deeper Probabilistic View of PDMs

When viewed through the lens of modern graphical modeling techniques and probabilistic principal component analysis, the classical formulation of PDMs can be re-written more formally as a probabilistic generative model over all landmarks

$\mathbf{x}$, where

$$
\begin{aligned}
p(\mathbf{x}) &= P(\mathbf{x}|\mathbf{z})\mathbf{P}(\mathbf{z}|\mathbf{q})\mathbf{P}(\mathbf{q}) \\
&= \int_{\mathbf{q}}\int_{\mathbf{z}} \mathcal{N}(\mathbf{x};\, s\mathbf{R}\mathbf{z}+\mathbf{t},\, \alpha\mathbf{I})\mathcal{N}(\mathbf{z};\, \mathbf{\Phi}\mathbf{q}+\boldsymbol{\mu},\, \sigma^{\mathbf{2}}\mathbf{I})\mathcal{N}(\mathbf{q};\, \mathbf{0},\, \mathbf{I})\, \mathbf{dq}\, \mathbf{dz}. \quad (40)
\end{aligned}
$$

Parameter $\mathbf{q}$ here acts as the prior of the $p(\mathbf{x})$ distribution with mean zero and unit covariance matrix. Given $\mathbf{q}$, the conditional $P(\mathbf{z}|\mathbf{q})$ gets its mean by applying the non-rigid deformation $\mathbf{\Phi}$ and adding the $\boldsymbol{\mu}$ term. The mean of the conditional $P(\mathbf{z}|\mathbf{q})$ sets the intermediate value in the mean of Eq. (39) before applying the rigid transformations. Finally, the mean of the conditional $P(\mathbf{x}|\mathbf{z})$ is measured by applying affine transformations of rotation $\mathbf{R}$, scaling s, and translation $\mathbf{t}$ on the variable $\mathbf{z}$. Having integrated out the variable $\mathbf{q}$, the distribution over all landmarks $\mathbf{x}$ can be simplified to

$$
p(\mathbf{x}) = \int_{\mathbf{z}} \mathcal{N}(\mathbf{x};\, s\mathbf{R}\mathbf{z}+\mathbf{t},\, \alpha\mathbf{I})\mathcal{N}(\mathbf{z};\, \boldsymbol{\mu},\, \mathbf{\Phi}\mathbf{\Phi}^{\mathbf{T}}+\sigma^{\mathbf{2}}\mathbf{I})\, \mathbf{dz}. \quad (41)
$$

Note that in this formulation the variable $\mathbf{z}$ represents the set of all landmarks in a transformed space where they are represented in a standardized coordinate space. That is different from their representation in $\mathbf{x}$ where each landmark is undergone an affine transformation and therefore can be observed in more volatile regions image. This difference guides us to a better representation in the $\mathbf{z}$ space to formulate the distribution of the key-points. Therefore, we transform the landmarks in each image from the $\mathbf{x}$ space to $\mathbf{z}$ space through a set of affine transformation, whose parameters are trained to be robust against translation, rotation, and scaling transformations. From this point forward, we discuss our model over the landmarks in the $\mathbf{z}$ space.

5.2. Probabilistic PCA based CLMs

Some of our previous work [24] has formulated a CLM in the following way; for a given image *I*, we wish to combine the outputs of the local classifiers for the keypoints with a spatial model of global keypoint configurations. The local classifiers acts as discriminative predictors for the keypoints, while the spatial model provides a prior on how the keypoints are distributed.

For each keypoint i a local SVM classifier is trained. At test time, the local classifier for keypoint i, generates a response image map $\mathbf{d}^i$, which is a 2D array with a probability prediction for each pixel position in the image being the keypoint i. The set $\mathbf{D} = \{\mathbf{d}^1, \mathbf{d}^2 \cdots, \mathbf{d}^n\}$ represents the set of all generated response images for a given image, where n is the number of keypoints. Figure 5 provides a visualization of the response image probability values. Note that the probabilities are scaled by a factor of 255 (8 bit gray-scale images).

Figure 5: Sample svm response maps generated by trained svms for an image. Each response map is generated by one local classifier trained for a particular keypoint.

Let $\mathbf{z}'_i \in \mathbf{d}^i$ indicate the coordinates of a gridpoint location on the 2D response image for keypoint i, then log of the score for the positive prediction of the local

classifier at this location can be defined as $s_i(\mathbf{z}'_i)$. We use a probabilistic PCA over the keypoints to model the global keypoint configurations and use the log of its Gaussian distribution with a factorized covariance matrix as it energy term and couple it with the log score of the local classifier predictions in a spatial interaction energy function as follows:

$$E(\boldsymbol{z}) = -\sum_{i=1}^{n} \sum_{\mathbf{z}'_i \in \mathbf{d}^i} s_i(\mathbf{z}'_i)\delta(\mathbf{z}_i - \mathbf{z}'_i) \tag{42}$$
$$+ \frac{1}{2}\left(\mathbf{z} - \boldsymbol{\mu}\right)^T \left(\boldsymbol{\Phi}\boldsymbol{\Phi}^{\mathbf{T}} + \sigma^{\mathbf{2}}\mathbf{I}\right)^{-\mathbf{1}} \left(\mathbf{z} - \boldsymbol{\mu}\right),$$

where $\mathbf{z} = [\mathbf{z}_1, \mathbf{z}_2, \cdots \mathbf{z}_n]^T$ gives the coordinates of n candidate locations for the keypoints whose energy is measured. The enegy function E is composed of two terms: The first term consists of the local response maps contribution where the first sum is over the keypoints and the second sum gets the log score $s_i(\mathbf{z}'_i)$ for each point $\mathbf{z}'_i$ in the set of all grid locations of the response image map $\mathbf{d}^i$. The terms δ is the dirac delta function whose output is one only if the grid location $\mathbf{z}'_i$ equals the queried location $\mathbf{z}_i$. The second term in Eq. (42) is log of the probabilistic PCA (PPCA) where $\boldsymbol{\mu}$ is simply the mean of the keypoints after RANSAC similarity registration. The terms $\boldsymbol{\Phi}$ and σ^2 in the covariance matrix of PPCA equal to:

$$\boldsymbol{\Phi} = \mathbf{U}_p(\boldsymbol{\Lambda}_p - \sigma^2\mathbf{I})^{\mathbf{1/2}}\mathbf{R}, \tag{43}$$

$$\sigma^2 = \frac{1}{n-p}\sum_{i=p+1}^{n} \lambda_i, \tag{44}$$

where $\mathbf{U}_p$ is a matrix of the p principle eingenvectors of the keypoints in $\mathbf{z}$ space, $\boldsymbol{\Lambda}_p$ is a diagonal matrix whose diagonal is composed of eigenvalues $\lambda_1, \ldots, \lambda_p$,

knowing that eigenvalue λ_j corresponds to the eigenvector in column j of the matrix $\mathbf{U_p}$, and $\mathbf{R}$ is an arbitrary orthogonal matrix. The term σ^2 is the average of the remaining eigenvalues which explain the least significant direction of variation in data. Note that the first and second terms in Eq. (42) correspond respectively to the first and second terms of Eq. (36).

To minimize E we perform a search over the candidate keypoint locations z. This is done by iterating over the keypoints, where in each iteration all key-points are fixed except one. The optimum value for that keypoint is found through a comprehnsive search over the energy of the local response map grid locations. The iterations continue until all keypoints converge.

5.3. Identity-Expression Factorized CLMs

Once the parameters of the identity expression model are learned using the EM procedure, as explained in Section (3.1), this model is used in conjunction with the local response classifiers derived from the previously trained SVMs. Combining the energy of the local classifier with the energy term of the Identity Expression Factorized model, we minimize the energy of the following function:

$$E(\boldsymbol{w}) = -\sum_{\mathbf{z}\in\boldsymbol{w}}\sum_{i=1}^{n}\sum_{\mathbf{z}'_i\in\mathbf{d}^i} s_i(\mathbf{z}'_i)\delta(\mathbf{z}_i-\mathbf{z}'_i) \tag{45}$$
$$+\left(\boldsymbol{w}-\boldsymbol{m}\right)^T\left(\boldsymbol{\psi}+\mathbf{A}\boldsymbol{\Phi}\mathbf{A}^T\right)^{-1}\left(\boldsymbol{w}-\boldsymbol{m}\right),$$

where $\boldsymbol{w}$ represents a set of candidate keypoints in all test images belonging to the same identity whose energy is measured. Note that as in the previous section $\mathbf{z}$ represents the set of all candidate keypoint locations on a single query image.

Assuming there are a total of $\boldsymbol{J}$ images in $\boldsymbol{w}$ representing different expressions of a given identity, the term $\boldsymbol{m}$ is simply $\boldsymbol{J}$ times concatenation of $\boldsymbol{\mu}$, which brings the mean of the keypoints to a new space of dimensionality $\boldsymbol{J} \times n$. Compared to the first term in Eq. (42), there is an added summation which iterates over all images belonging to the same identity. The second term in Eq. (45), is equivalent to the log of the marginal distribution $P(\boldsymbol{w})$ of the identity expression factorized model, explained in Section 3, which gets the probability of the joint set of points in all images of the same identity being the key-points. The term $\boldsymbol{\Psi}$ is constructed as a diagonal matrix by concatenating J_i times the diagonal of Σ and the matrix definitions of $\mathbf{A}$ and $\boldsymbol{\Phi}$ are given respectively in Eq. (35) and Eq. (34).

The optimization procedure is done as follows; for each key-point $i \in n$ in each image $j \in \boldsymbol{J}$ a SVM response map $\mathbf{d}^i$ is generated. In each response map the position with the highest probability is chosen as the key-point location. The set of selected locations give the initial value of $\boldsymbol{w}$. Then, using Iterated Conditional Modes (ICM) [5], the energy of (45) is minimized by updating the position of each keypoint to it's minimum energy position while having all other keypoints fixed, iterating over all keypoints multiple times until no further updates are possible.

5.4. Keypoint Localization Experiments with the MultiPIE Dataset and IE-CLMs

The CMU MultiPie face dataset [23] captures a subset of expressions, such as surprise, smile, neutral, squint, disgust, and scream, of 337 identities. The images are taken in four sessions and in each session a subset of expressions is captured. Therefore, not all expressions are registered per identity. The availability of multiple expressions per identity provides the right dataset to evaluate our identity expression model. We perform the same experiment as in [56] and compare our model with other models on the frontal face images with 68 keypoints. Following

Figure 6: Sample of keypoint localization for multiple expressions of different identities by IE-CLM model.

the same split of the dataset as in [56], the first 300 images are used for training and the next 300 images are used for testing. Figure 7 shows the main directions of variation for the identity and expression features captured by our model on MulitPie training set. The directions of variation for expression captures how face changes in between different emotions. On the other hand, the directions of variation for identity mainly deals with different face sizes as well as some 2D in-plane rotation which is due to the fact that the faces are not completely frontal in some of the expressions of an identity.

Using the training set of MultiPie, we train a binary SVM classifier on HOG feature representation for each keypoint. Indeed, for training the SVM classifier of each keypoint, positive patches are generated from the training set centered at the keypoint and negative patches are generated within a bounding box around the keypoint. We compare our model on MultiPie test-set with the independent model

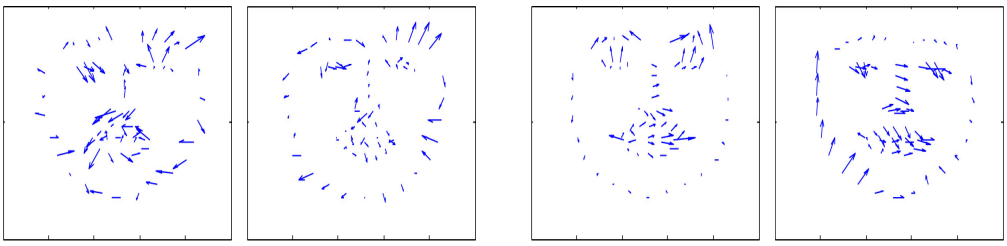

1st (left) and 2nd (right) directions for id. 1st (left) and 2nd (right) directions for exp.

Figure 7: The main directions of variation across identities and expressions

of Zhu and Ramanan [56] and other models reported in that paper as well as the more recent model of Hasan *et al.* [24]. The error is the average difference in euclidean space between the true pixel locations and the predicted ones normalized by the face size, which is the average of height and width. The results are reported in Table 5. Our model catches up with Zhu's model at 0.06 of face fraction size, though at face fraction size of 0.05 the difference is marginal at 0.016. The values of the CLM model of Saragih *et al.* [39] are reported from [56] where the latter source evaluates the previously trained CLM model on MultiPie test-set. However, it is not clear on which data the CLM model was trained on and which feature representation was used for training them. In order to provide a CLM model trained on precisely the same training set and using precisely the same features and local classifiers, we use the PPCA-CLMs described in Section 5.2, which are trained on HOG features. This allows us to establish a more highly controlled baseline for the PCA vs IE-CLM comparison. The performance of our PPCA-CLM is different from the PCA-CLM of Saragih *et al.*. These differences could be due to the difference in the training set used, or due to the difference in the feature representation. Other variations could arise from our slightly different

underlying CLM formulation or our use of ICM as the inference procedure. We see that for small face fraction size accuracy the difference is significant; however, at the 0.06 face fraction threshold, the performance of the two PCA based CLM methods are comparable.

Turning to the IE-CLM we see that it lags behind Saragih's CLM only at the very small face fraction accuracy level. We speculate that the lower performance of IE-CLM in this range is due to the difference in the training data for the local classifiers. The IE-CLM model catches up rapidly to both Saragih's CLM and the Multi-AAM [29] models from the 0.04 face fraction level forward, and yields predictions for 100% of the points with an error of less than 0.06 of the face size. This level of performance is on par with the state of the art methods of [56] and [24]. Figures 8 and 9 illustrate our model's performance on different expressions of two sample identities in the test-set. Our model exploits the fact that all these images have the same identity and thereupon, the same identity representation is used when the model searches for the optimal keypoints. Figure 6 shows some sample keypoint localization by our model on different identities in the test set, which includes data of different expression, ethnicity, age, gender, illumination, and face masks such as trimming style and eye-glasses usage.

In terms of running time, the PPCA-CLM model takes 32 seconds for the local response image generation and 6 seconds for the optimization procedure. However, we have use none of the obvious acceleration techniques that could be applied to speed the local response map generation procedure, such as GPU acceleration or optimizing operations that are convolutational in nature. The IE-CLM model's time changes linearly based on the number of images of the identity being evaluated. If intended for an interactive or real-time procedure where the images

Figure 8: Labeled images by Identity Expression Factorized CLM model for different expressions of identity sample 1.

are to be processed sequentially, the timing could be made much more comparable to a PCA-CLM for which various real-time implementations are available. To accelerate the IE-CLM computations one could adapt the joint inference procedure over all images in a batch so as to *incrementally* use estimated factors from images in previous time steps. This could be formulated as a form of incremental inference or probability propagation, an approach which is fairly well understood in the literature.

6. Conclusions and Future Work

We have shown that identity-expression disentanglement and various forms of identity normalization can be used to improve the performance of supervised learning approaches to performance-driven animation and expression recognition tasks. Detailed facial expression analysis techniques are often limited by the need for subject-specific models, which limits the quantity of labeled data that can be

Figure 9: Labeled images by Identity Expression Factorized CLM model for different expressions of identity sample 2.

brought to bear. In our work here, we have detailed how approaches that work well for a single subject can be adapted and their performance improved for unseen identities. Our approach is based on disentangling identity and expression with data that was been weakly labeled according to identity. The increase in performance across many tasks suggests that using simpler tools for both performance-driven animation and automatic expression recognition may be sufficient if identity labels can be used as an additional source of information.

In future work it might be interesting to extend this type of approach to be able to deal with more complex data, including data exhibiting greater pose variation. For example an explicit 3D shape model such as the formulation proposed for CLMs in the work of Jeni et al. [28] could be combined with the identity-expression factorization approach used here.

We have shown how the identity-expression factorization approach can be integrated into a CLM for keypoint localization. We believe there are a number of

possible directions that could lead to further improvements. Firstly, we believe that replacing the binary SVM classifier with a multi-class SVM might yield increased performance with minimal changes required to the underlying model formulation. However, once formulated in this way, it would also be more straightforward to train the model to make structured predictions using a fully discriminative training procedure, for example, along the lines of [56]. Another direction is to replace the Iterated Conditional Modes (ICM) energy minimization step with a temperature based annealing procedure. The ICM algorithm is known to be quite fast at finding local minima, but it may be the case that better solutions could be found with a stochastic annealing procedure.

The use of multi-scale intensity response information in conjunction with alternative optimization techniques could be a promising direction for future research. It is also likely possible to speed up response image generation and energy minimization when multiple images are processed per identity through various means.

Finally, the local SVM classifiers and hand engineered features that we have used here could be replaced with a convolutional neural network (CNN), such as the architecture proposed in [25]. A CNN approach was not feasible in the context of our primary goal of animation control at the time this work was performed since CNNs typically require a considerable amount of training data to avoid overfitting and our face camera dataset is quite small compared to the data sets that have been used where CNNs have yielded state of the art results. However, for other application settings and datasets research combining identity-expression disentanglement techniques with CNNs may be a promising direction of future research.

Acknowledgements

We thank Ubisoft for both financial support and for providing the helmet camera video data used for our high quality animation control experiments. We also thank the Natural Sciences and Engineering Research Council of Canada (NSERC) for financial support under the CRD and Discovery grant programs and the Fonds de recherche du Quebec - Nature et technologies (FRQNT) for a Doctoral research scholarships (B2) grant to SH. We thank Yoshua Bengio for his comments and suggestions regarding this work.

References

[1] http://www.face.com.

[2] F. R. Bach and M. I. Jordan. A probabilistic interpretation of canonical correlation analysis. Technical Report 688, Department of Statistics, University of California, Berkeley, 2005.

[3] M. S. Bartlett, J. C. Hager, P. Ekman, and T. J. Sejnowski. Measuring facial expressions by computer image analysis. *Psychophysiology*, 36(2):253–263, Mar. 1999.

[4] T. Beeler, B. Bickel, P. Beardsley, B. Sumner, and M. Gross. High-quality single-shot capture of facial geometry. *ACM Trans. Graph.*, 29(4):40:1—-40:9, July 2010.

[5] J. Besag. On the statistical analysis of dirty pictures. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 259–302, 1986.

[6] V. Blanz and T. Vetter. A morphable model for the synthesis of 3d faces. pages 187–194. ACM Press, 1999.

[7] Y. Chang, M. Vieira, M. Turk, and L. Velho. Automatic 3D Facial Expression Analysis in Videos. In W. Zhao, S. Gong, and X. Tang, editors, *Analysis and Modelling of Faces and Gestures SE - 23*, volume 3723 of *Lecture Notes in Computer Science*, pages 293–307. Springer Berlin Heidelberg, 2005.

[8] F. Cheng, J. Yu, and H. Xiong. Facial expression recognition in JAFFE dataset based on gaussian process classification. *Neural Networks, IEEE Transactions on*, 21(10):1685–1690, 2010.

[9] J. F. Cohn and M. A. Sayette. Spontaneous facial expression in a small group can be automatically measured: An initial demonstration. *Behavior Research Methods*, 42(4):1079–1086, Nov. 2010.

[10] T. Cootes, C. Taylor, D. Cooper, and J. Graham. Active Shape Models-Their Training and Application. *Computer Vision and Image Understanding*, 61(1):38–59, Jan. 1995.

[11] D. Cristinacce and T. F. Cootes. Feature detection and tracking with constrained local models. In *BMVC*, volume 2, page 6, 2006.

[12] A. Del Bue, J. a. Xavier, L. Agapito, and M. Paladini. Bilinear modeling via augmented Lagrange multipliers (BALM). *IEEE transactions on pattern analysis and machine intelligence*, 34(8):1496–508, Aug. 2012.

[13] Y. Du and X. Lin. Emotional facial expression model building. *Pattern Recognition Letters*, 24(16):2923–2934, Dec. 2003.

[14] G. Edwards, A. Lanitis, C. Taylor, and T. Cootes. Statistical models of face images . improving specificity. *Image and Vision Computing*, 16(3):203–211, Mar. 1998.

[15] P. Ekman et al. An argument for basic emotions. *Cognition and emotion*, 6(3/4):169200, 1992.

[16] I. Essa, S. Basu, T. Darrell, and A. Pentland. Modeling, tracking and inter-

active animation of faces and heads//using input from video. In *Computer Animation '96. Proceedings*, pages 68–79, 1996.

[17] M. R. Everingham, J. Sivic, and A. Zisserman. Hello! my name is... buffy" – automatic naming of characters in tv video. In *Proc. BMVC*, pages 92.1–92.10, 2006. doi:10.5244/C.20.92.

[18] P. F. Felzenszwalb, R. B. Girshick, D. McAllester, and D. Ramanan. Object detection with discriminatively trained part-based models. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 32(9):1627–1645, 2010.

[19] W. Freeman and J. Tenenbaum. Learning bilinear models for two-factor problems in vision. pages 554–560, 1997.

[20] B. J. Frey and J. Nebojsa. Estimating mixture models of images and inferring spatial transformations using the EM algorithm. pages 416–422, Ft. Collins, CO, USA, 1999.

[21] K. Goudeaux, T. Chen, S. Wang, and J. Liu. Principal component analysis for facial animation. In *Acoustics, Speech, and Signal Processing, 2001. Proceedings. (ICASSP '01). 2001 IEEE International Conference on*, volume 3, pages 1501–1504, 2001.

[22] R. Gross, I. Matthews, and S. Baker. Generic vs. person specific active appearance models. *Image and Vision Computing*, 23(12):1080–1093, Nov. 2005.

[23] R. Gross, I. Matthews, J. Cohn, T. Kanade, and S. Baker. Multi-pie. *Image Vision Comput.*, 28(5):807–813, May 2010.

[24] M. K. Hasan, C. Pal, and S. Moalem. Localizing facial keypoints with global descriptor search, neighbour alignment and locally linear models. In *Proceedings of the 2013 IEEE International Conference on Computer Vision Workshops*, ICCVW '13, pages 362–369, Washington, DC, USA, 2013.

IEEE Computer Society.

[25] S. Honari, J. Yosinski, P. Vincent, and C. J. Pal. Recombinator networks: Learning coarse-to-fine feature aggregation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016.

[26] G. Huang, V. Jain, and E. Learned-Miller. Unsupervised joint alignment of complex images. In *Computer Vision, 2007. ICCV 2007. IEEE 11th International Conference on*, pages 1–8, 2007.

[27] R. G. Jason Saragih. Learning AAM fitting through simulation. *Pattern Recognition*, 42(11):2628—-2636, 2009.

[28] L. A. Jeni, A. Lőrincz, T. Nagy, Z. Palotai, J. Sebők, Z. Szabó, and D. Takács. 3d shape estimation in video sequences provides high precision evaluation of facial expressions. *Image and Vision Computing*, 30(10):785–795, 2012.

[29] D.-J. Kroon. Active shape model and active appearance model. *http://www.mathworks.com/matlabcentral/fileexchange/26706*.

[30] A. Lanitis, C. J. Taylor, and T. F. Cootes. Automatic interpretation and coding of face images using flexible models. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 19(7):743–756, July 1997.

[31] P. Liao, L. Shen, Y. Chen, and S. Liu. Unified model in identity subspace for face recognition. *Journal of Computer Science and Technology*, 19(5):684–690, 2004.

[32] P. Lucey, J. F. Cohn, T. Kanade, J. Saragih, Z. Ambadar, and I. Matthews. The Extended Cohn-Kanade Dataset (CK+): A complete dataset for action unit and emotion-specified expression. In *2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition - Workshops*, pages 94–101. IEEE, June 2010.

[33] M. Lyons, S. Akamatsu, M. Kamachi, and J. Gyoba. Coding facial expres-

sions with gabor wavelets. In *Automatic Face and Gesture Recognition, 1998. Proceedings. Third IEEE International Conference on*, pages 200–205. IEEE, 1998.

[34] I. Matthews and S. Baker. Active appearance models revisited. *International Journal of Computer Vision*, 60:135–164, Nov. 2004.

[35] A. Pentland, B. Moghaddam, and T. Starner. View-based and modular eigenspaces for face recognition. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition CVPR-94*, pages 84–91. IEEE Comput. Soc. Press, 1994.

[36] F. Pighin, R. Szeliski, and D. Salesin. Resynthesizing facial animation through 3D model-based tracking. In *Computer Vision, 1999. The Proceedings of the Seventh IEEE International Conference on*, volume 1, pages 143–150 vol.1, 1999.

[37] S. Prince, P. Li, Y. Fu, U. Mohammed, and J. Elder. Probabilistic models for inference about identity. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, PP(99):1, 2011.

[38] G. Sandbach, S. Zafeiriou, M. Pantic, and L. Yin. Static and dynamic 3D facial expression recognition: A comprehensive survey. *Image and Vision Computing*, 30(10):683–697, Oct. 2012.

[39] J. M. Saragih, S. Lucey, and J. F. Cohn. Deformable Model Fitting by Regularized Landmark Mean-Shift. *International Journal of Computer Vision*, 91(2):200–215, Sept. 2010.

[40] D. Sibbing, M. Habbecke, and L. Kobbelt. Markerless reconstruction of dynamic facial expressions. In *Computer Vision Workshops (ICCV Workshops), 2009 IEEE 12th International Conference on*, pages 1778–1785, 2009.

[41] Y. Sun, X. Wang, and X. Tang. Deep convolutional network cascade for

facial point detection. In *Computer Vision and Pattern Recognition (CVPR), 2013 IEEE Conference on*, pages 3476–3483. IEEE, 2013.

[42] M. Suwa, N. Sugie, and K. Fujimora. A preliminary note on pattern recognition of human emotional expression. In *International Joint Conference on Pattern Recognition*, pages 408–410, 1978.

[43] Y. Tian, T. Kanade, and J. F. Cohn. Facial expression recognition. In S. Z. Li and A. K. Jain, editors, *Handbook of Face Recognition*, pages 487–519. Springer London, London, 2011.

[44] M. Tistarelli and M. S. Nixon, editors. *Advances in Biometrics*, volume 5558 of *Lecture Notes in Computer Science*. Springer Berlin Heidelberg, Berlin, Heidelberg, 2009.

[45] M. Valstar, I. Patras, and M. Pantic. Facial action unit detection using probabilistic actively learned support vector machines on tracked facial point data. In *Computer Vision and Pattern Recognition - Workshops, 2005. CVPR Workshops. IEEE Computer Society Conference on*, page 76, 2005.

[46] M. F. Valstar, B. Jiang, M. Mehu, M. Pantic, and K. Scherer. The first facial expression recognition and analysis challenge. In *2011 IEEE International Conference on Automatic Face & Gesture Recognition and Workshops (FG 2011)*, pages 921–926. IEEE, Mar. 2011.

[47] L. van der Maaten and E. Hendriks. Capturing appearance variation in active appearance models. In *2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition - Workshops*, pages 34–41. IEEE, June 2010.

[48] M. A. O. Vasilescu and D. Terzopoulos. Multilinear analysis of image ensembles: TensorFaces. In A. Heyden, G. Sparr, M. Nielsen, and P. Jo-

hansen, editors, *Computer Vision ECCV 2002*, volume 2350, pages 447–460. Springer Berlin Heidelberg, Berlin, Heidelberg, 2002.

[49] L. Vezzaro. Icaam - inverse compositional active appearance models, November 2012. `http://bit.ly/173xL6b`.

[50] D. Vlasic, M. Brand, H. Pfister, and J. Popovi. Face transfer with multilinear models. *ACM Trans. Graph.*, 24(3):426433, July 2005.

[51] Y. Wang, X. Huang, C. Lee, S. Zhang, Z. Li, D. Samaras, D. Metaxas, A. Elgammal, and P. Huang. High resolution acquisition, learning and transfer of dynamic 3D facial expressions. *Computer Graphics Forum*, 23(3):677–686, Sept. 2004.

[52] L. Williams. Performance-driven facial animation. *SIGGRAPH Comput. Graph.*, 24(4):235242, Sept. 1990.

[53] L. Yin, X. Chen, Y. Sun, T. Worm, and M. Reale. A high-resolution 3D dynamic facial expression database, 2008.

[54] L. Zhang, N. Snavely, B. Curless, and S. M. Seitz. Spacetime faces: High-Resolution capture for˜Modeling and animation. In Z. Deng and U. Neumann, editors, *Data-Driven 3D Facial Animation*, pages 248–276. Springer London, London, 2007.

[55] C. Zhou and X. Lin. Facial expressional image synthesis controlled by emotional parameters. *Pattern Recognition Letters*, 26(16):2611–2627, Dec. 2005.

[56] X. Zhu and D. Ramanan. Face detection, pose estimation, and landmark localization in the wild. In *Computer Vision and Pattern Recognition (CVPR), 2012 IEEE Conference on*, pages 2879–2886. IEEE, 2012.

Table 1: Summary of experiments described in this paper, numbers of the section describing each experiment are given in parenthesis.

JAFFE		
Emotion Recognition		(3.2.1)
predicts:	Emotion labels	
using:	Images	
Animation Control		(3.2.3)
predicts:	Bone position parameters	
using:	Images	
Extended Cohn-Kanade		
Facial Action Unit (AU) Detection		(3.2.2)
predicts:	AU labels	
using:	Point locations	
	Shape-normalized images	
	Combined point locations and shape-normalized images	
Emotion Recognition		(3.2.2)
predicts:	Emotion labels	
using:	Point locations	
	Shape-normalized images	
	Combined point locations and shape-normalized images	
Key-Point Localization (IE-AAMs)		(4.3)
predicts:	Key-point locations	
using:	Images	
CMU Multi-PIE		
Key-Point Localization (IE-CLMs)		(5.4)
predicts:	Key-point locations	
using:	Images	
Animation Control Studio Data		
Animation Control		(3.2.4)
predicts:	Bone position parameters	
using:	Shape-normalized images	

Table 2: Accuracy for JAFFE emotion recognition in percentage and Mean Squared Error for bone position recovery experiments for JAFFE and Studio Motion Capture data, calculated per bone position, which lie in $[-1, 1]$.

	JAFFE				Studio Data	
	Emotion Recognition		Bone Position Recovery		Bone Position Recovery	
	Accuracy	SE	MSE	SE	MSE	SE
No Factor Analysis	53.13%	3.39%	1.7526	0.2702	1.5809	0.2300
PCA 100 dimensions	57.08%	6.57%	1.7526	0.2702	0.0786	0.0120
PCA 30 dimensions	56.13%	5.64%	0.1223	0.0121	0.1007	0.0123
Our Method	**72.71%**	1.83%	**0.0851**	0.0077	**0.0231**	0.0021

Table 3: CK+: AUC Results and estimated standard errors of the AU experiment

CK+ Facial Action Unit Recognition

		Baseline [32]			Identity Normalized		
AU	N	SPTS	CAPP	SPTS+CAPP	SPTS	CAPP	SPTS+CAPP
1	173	94.1 ± 1.8	91.3 ± 2.1	96.9 ± 1.3	97.7 ± 2.8	96.33 ± 3.2	$\mathbf{99.1 \pm 1.2}$
2	116	97.1 ± 1.5	95.6 ± 1.9	97.9 ± 1.3	97.2 ± 2.4	97.82 ± 1.0	$\mathbf{98.5 \pm 0.9}$
4	191	85.9 ± 2.5	83.5 ± 2.7	91.0 ± 2.1	89.6 ± 9.3	91.72 ± 7.7	$\mathbf{94.3 \pm 5.6}$
5	102	95.1 ± 2.1	96.6 ± 1.8	97.8 ± 1.5	96.1 ± 2.1	97.52 ± 2.6	$\mathbf{98.0 \pm 1.6}$
6	122	91.7 ± 2.5	94.0 ± 2.2	95.8 ± 1.8	96.8 ± 3.3	95.37 ± 4.1	$\mathbf{97.3 \pm 2.8}$
7	119	78.4 ± 3.8	85.8 ± 3.2	89.2 ± 2.9	89.7 ± 7.4	92.95 ± 7.1	$\mathbf{94.7 \pm 5.9}$
9	74	97.7 ± 1.7	99.3 ± 1.0	99.6 ± 0.7	99.2 ± 0.8	$\mathbf{99.71 \pm 0.4}$	98.2 ± 0.2
11	33	72.5 ± 7.8	82.0 ± 6.7	85.2 ± 6.2	81.1 ± 4.1	83.15 ± 4.2	$\mathbf{91.6 \pm 3.6}$
12	111	91.0 ± 2.7	96.0 ± 1.9	96.3 ± 1.8	96.9 ± 2.2	$\mathbf{97.68 \pm 2.2}$	97.0 ± 2.0
15	89	79.6 ± 4.3	88.3 ± 3.4	89.9 ± 3.2	90.2 ± 4.2	96.42 ± 2.0	$\mathbf{96.8 \pm 2.5}$
17	196	84.4 ± 2.6	90.4 ± 2.1	93.3 ± 1.8	92.1 ± 6.7	95.48 ± 4.2	$\mathbf{96.3 \pm 3.6}$
20	77	91.0 ± 3.3	93.0 ± 2.9	94.7 ± 2.6	95.9 ± 3.4	94.62 ± 2.4	$\mathbf{96.7 \pm 1.7}$
23	59	91.1 ± 3.7	87.6 ± 4.3	92.2 ± 3.5	91.7 ± 3.7	93.77 ± 2.9	$\mathbf{96.5 \pm 2.8}$
24	57	83.3 ± 4.9	90.4 ± 3.9	91.3 ± 3.7	88.9 ± 4.5	92.91 ± 3.4	$\mathbf{94.2 \pm 3.4}$
25	287	97.1 ± 1.0	94.0 ± 1.4	97.5 ± 0.9	97.6 ± 2.6	96.83 ± 3.6	$\mathbf{98.1 \pm 1.9}$
26	48	75.0 ± 6.3	77.6 ± 6.0	80.3 ± 5.7	77.9 ± 8.1	83.61 ± 7.4	$\mathbf{86.6 \pm 6.8}$
27	81	99.7 ± 0.7	98.6 ± 1.3	99.8 ± 0.5	99.9 ± 0.2	99.67 ± 0.5	$\mathbf{99.9 \pm 0.2}$
AVG		90.0 ± 2.5	91.4 ± 2.4	94.5 ± 2.0	92.8 ± 4.0	94.45 ± 3.5	$\mathbf{96.1 \pm 2.7}$

Table 4: Comparison of confusion matrices of emotion detection for the combined landmark (SPTS) and shape-normalised image (CAPP) features before and after identity normalisation. The average accuracy for all predicted emotions using the state of the art method in [32] is 83.27% (top table), using our method yields 95.21% (bottom table), a substantial improvement of 11.9%.

CK+ Emotion Recognition

Baseline SPTS+CAPP [32]

	An	Di	Fe	Ha	Sa	Su	Co
An	**75.0**	7.5	5.0	0.0	5.0	2.5	5.0
Di	5.3	**94.7**	0.0	0.0	0.0	0.0	0.0
Fe	4.4	0.0	**65.2**	8.7	0.0	13.0	8.7
Ha	0.0	0.0	0.0	**100.0**	0.0	0.0	0.0
Sa	12.0	4.0	4.0	0.0	**68.0**	4.0	8.0
Su	0.0	0.0	0.0	0.0	4.0	**96.0**	0.0
Co	3.1	3.1	0.0	6.3	3.1	0.0	**84.4**

Identity Normalized SPTS+CAPP

	An	Di	Fe	Ha	Sa	Su	Co
An	**95.4**	0.0	0.0	0.0	0.0	0.0	4.7
Di	1.9	**94.2**	0.0	0.0	1.9	0.0	1.9
Fe	0.0	0.0	**84.2**	0.0	0.0	5.3	10.5
Ha	0.0	0.0	1.6	**96.9**	0.0	0.0	1.6
Sa	0.0	0.0	5.0	0.0	**90.0**	0.0	5.0
Su	0.0	0.0	0.0	0.0	0.0	**98.8**	1.2
Co	7.7	0.0	0.0	0.0	0.0	0.0	**92.3**

Table 5: Percentage of faces with an average localization error less than the given fraction of face size on MultiPie dataset.

Fraction of face size	0.03	0.04	0.05	0.06
Zhu and Ramanan (independent model) [56]	0.86	0.97	1.00	1.00
Hasan *et al.* [24]	0.78	0.94	1.00	1.00
Oxford [17]	0.28	0.77	0.94	0.98
Star Model [18]	0.29	0.80	0.92	0.92
Multi-AAM [29]	0.64	0.87	0.91	0.92
CLM of Saragih *et al.* [39]	0.68	0.85	0.90	0.93
Face.com [1]	0.31	0.61	0.79	0.87
Our PPCA-CLM	0.13	0.54	0.79	0.92
Our IE-CLM	0.52	0.89	0.98	1.00